

The parameterized space complexity of embedding along a path

Hubie Chen¹ and Moritz Müller²

¹ University of the Basque Country (UPV/EHU), E-20018 San Sebastián, Spain, *and*
IKERBASQUE, Basque Foundation for Science, E-48011 Bilbao, Spain,

hubie.chen@ehu.es

² Kurt Gödel Research Center, University of Vienna, Austria ,

moritz.mueller@univie.ac.at

Abstract

The embedding problem is to decide, given an ordered pair of structures, whether or not there is an injective homomorphism from the first structure to the second. We study this problem using an established perspective in parameterized complexity theory: the universe size of the first structure is taken to be the parameter, and we define the embedding problem relative to a class \mathcal{A} of structures to be the restricted version of the general problem where the first structure must come from \mathcal{A} . We initiate a systematic complexity study of this problem family, by considering classes whose structures are what we call rooted path structures; these structures have paths as Gaifman graphs. Our main theorem is a dichotomy theorem on classes of rooted path structures.

1 Introduction

The *embedding problem* is to decide, given a pair (\mathbf{A}, \mathbf{B}) of structures, whether or not there is an *embedding*—an injective homomorphism—from \mathbf{A} to \mathbf{B} . Intuitively, the embedding problem asks whether one can identify a particular type of pattern, specified by the first structure \mathbf{A} , in the second structure \mathbf{B} . Of course, this is a problem of a fundamental nature. Indeed, a number of well-established computational problems can be viewed as cases of this problem. Examples include the problems CLIQUE, CYCLE and PATH which ask, given an undirected graph \mathbf{G} and a natural number k , whether \mathbf{G} contains a size k clique, length k cycle or length k path, respectively.

In parameterized complexity theory, the embedding problem is typically studied by taking the universe size of the first structure \mathbf{A} as the parameter; this is the perspective and parameterization that we use here. In the examples above, this corresponds to the standard parameterization by the value k , and yields the famous parameterized problems p -CLIQUE, p -CYCLE and p -PATH. While this parameterized problem is in general intractable (p -CLIQUE is W[1]-complete), it has been fruitful to consider the following family of restricted versions of the problem: for each class \mathcal{A} of structures, p -EMB(\mathcal{A}) is the parameterized embedding problem where the first structure \mathbf{A} must come from \mathcal{A} . Throughout, we assume that each class of structures under discussion is on a shared finite vocabulary. In a now-famous result [1], it was established that the problem p -EMB(\mathcal{A}) is fixed-parameter tractable when \mathcal{A} has bounded treewidth. The algorithmic technique introduced there, called *color coding*, in fact can be viewed as providing a Turing reduction from the embedding problem p -EMB(\mathcal{A}) to the homomorphism problem p -HOM(\mathcal{A}^*). Here, the problem p -HOM(\cdot) is

defined analogously to $p\text{-EMB}(\cdot)$, but asks merely for a homomorphism (as opposed to an injective homomorphism); the class \mathcal{A}^* is obtained from \mathcal{A} by replacing each structure \mathbf{A} in \mathcal{A} with the structure \mathbf{A}^* , which is the structure \mathbf{A} but expanded so that there is, for each element a of \mathbf{A} , a unary relation symbol U_a interpreted as $\{a\}$.

The family $p\text{-HOM}(\mathcal{A})$ of homomorphism problems is well-understood. A classification of these problems up to parameterized logarithmic space reduction is known [3], which shows that each problem $p\text{-HOM}(\mathcal{A})$ is either in para-L (parameterized logarithmic space), PATH-complete, TREE-complete, or W[1]-complete. The complexity classes PATH and TREE are subclasses of FPT, and indeed the inclusions $\text{para-L} \subseteq \text{PATH} \subseteq \text{TREE} \subseteq \text{FPT}$ are known (further discussion of these classes can be found in [3]). It has been shown that para-L and PATH are not equal under the assumption that Savitch's classical simulation cannot be improved [4].

In contrast, the family $p\text{-EMB}(\mathcal{A})$ of embedding problems seems quite enigmatic. While it has been conjectured that the problem $p\text{-EMB}(\mathcal{A})$ is W[1]-hard when \mathcal{A} does not have bounded treewidth (see for example [7, p.355]), this research issue is (to our knowledge) wide open. Indeed, only recently was the complexity of the prominent problem BICLIQUE resolved as W[1]-complete [8]; this can be defined as the particular problem $p\text{-EMB}(\mathcal{A})$ where \mathcal{A} is the class of complete bipartite graphs. Let us mention that, concerning our examples, it is known that $p\text{-PATH}$ is in para-L [4] and that $p\text{-CYCLE}$ is PATH-complete [3].


Contributions The motivation behind the present work was to initiate a systematic study of the $p\text{-EMB}(\mathcal{A})$ family of problems, in hopes of eventually obtaining classification results of the form known for the homomorphism problem. We here focus on classes \mathcal{A} of *rooted path structures*. These are structures whose Gaifman graph is a path and in addition are *rooted* in the sense that one of its endpoints is the sole element of a relation. While this implies that $p\text{-EMB}(\mathcal{A})$ is in PATH (and hence in FPT), the suggestion here is to first obtain a thorough understanding of the problem family with respect to small complexity classes, and then attempt to scale up this understanding.

As examples of our findings, consider the following three classes of structures.


1. The class of rooted alternating paths.



2. The class derived from rooted alternating paths by subdividing each edge.



3. The class derived from rooted alternating paths by adding one final non-alternating edge:




It follows from our results that, with respect to the embedding problem, the first class is in para-L, whereas the second and third one are PATH-complete. The complexity of these classes can be derived from a dichotomy theorem (Theorem 5.1) that characterizes classes of rooted path structures where each structure is an oriented path with a root.

Our main theorem is a dichotomy theorem which describes the complexity of the problem $p\text{-EMB}(\mathcal{A})$ for each class of rooted path structures (Theorem 3.2). However, we do not succeed in obtaining a para-L versus PATH-complete dichotomy, as in the previously mentioned dichotomy theorem. Instead, for each such problem $p\text{-EMB}(\mathcal{A})$, we either show it to be PATH-complete or we exhibit a parameterized logarithmic

space algorithm that solves the problem with oracle access to a problem which we call the *long-short path problem*. This algorithmic result is based on color coding and Reingold's algorithm.

A number of remarks are in order. First, if one shows that the long-short path problem is in para-L, our positive complexity result can immediately be improved to containment in para-L. Second, with respect to the problem family considered, the long-short path problem is both unavoidable and occurs naturally in the family, in the following precise sense: there exists a class of rooted path structures \mathcal{A} such that $p\text{-EMB}(\mathcal{A})$ is equivalent, under parameterized logarithmic space Turing reduction, to the long-short path problem (Theorem 6.1). Hence, one necessarily needs to resolve the complexity of the long-short path problem in order to describe all problems in the studied problem family up to parameterized logarithmic space Turing reduction. Third, independently of what the complexity of the long-short path problem turns out to be, we believe that the present work makes a contribution in identifying and isolating the long-short path problem as the hardest of the embedding problems in our family that are of unknown complexity. This identification can indeed be conceived of as a form of completeness result. Our view is that settling the complexity of this concrete problem is a challenge to known techniques, and thus that a deeper understanding of this problem could mark healthy progress in the understanding of parameterized logarithmic space.

To get some feeling for the difficulty, we encourage the reader to ponder whether the embedding problem associated with the following class of rooted path structures is in para-L or PATH-complete. The class consists in undirected paths prolonged by alternating paths. In a picture:



This gives the maybe the simplest (and most annoying) example of an embedding problem for rooted path structures which we conjecture to belong to para-L but are only able to reduce to the long-short path problem.

To close this introduction, let us make the following observations. As mentioned, it is known that when the class \mathcal{A} has bounded treewidth, the problem $p\text{-EMB}(\mathcal{A})$ is in FPT; it has been conjectured that the problem $p\text{-EMB}(\mathcal{A})$ is W[1]-hard otherwise. This conjecture thus suggests that one need only look at the Gaifman graphs of the structures in a class \mathcal{A} to determine whether or not $p\text{-EMB}(\mathcal{A})$ is in FPT. Under the assumption that the complexity degrees dealt with in this article are pairwise distinct, our results contrast sharply with this suggestion: we only consider structures with path Gaifman graphs, but show that within the realm of such structures, dichotomies occur, and hence the Gaifman graph does not carry the information needed to determine the complexity of $p\text{-EMB}(\mathcal{A})$. Moreover, the aforementioned conjecture implies that if $p\text{-EMB}(\mathcal{A})$ is in FPT at all, then it is in FPT via using color coding to reduce to $p\text{-HOM}(\mathcal{A}^*)$. Our positive complexity results go strictly beyond this paradigm of reducing to $p\text{-HOM}(\mathcal{A}^*)$ because $p\text{-HOM}(\mathcal{A}^*)$ is always PATH-complete for any infinite class \mathcal{A} of rooted path structures over the same finite vocabulary (this follows from [3]).

2 Preliminaries

2.1 Structures and logic

A (*relational*) *vocabulary* is a finite set τ of relation symbols; every $R \in \tau$ has an associated *arity* $r \in \mathbb{N}$. Recall τ -*formulas* are built from atomic τ -formulas by means of \wedge , \neg and $\exists x$, and an *atomic* τ -formula has the form $R(x_1, \dots, x_r)$ or $x_1 = x_2$ where the x_i are variables and R is an r -ary relation symbol from τ . The notation $\varphi(x_1, \dots, x_r)$ means that the free variables of the τ -formula φ are among x_1, \dots, x_r .

A τ -structure \mathbf{A} consists of a non-empty *universe* A and for every r -ary relation symbol $R \in \tau$ a relation $R^{\mathbf{A}} \subseteq A^r$. We only consider structures with finite universes. A (*induced*) *substructure* of \mathbf{A} is a τ -structure \mathbf{B} with $B \subseteq A$ and $R^{\mathbf{B}} = R^{\mathbf{A}} \cap B^r$ for every r -ary $R \in \tau$.

If $\varphi(x_1, \dots, x_r)$ is a τ -formula, $\mathbf{A} \models \varphi(a_1, \dots, a_r)$ means that the tuple $(a_1, \dots, a_r) \in A^r$ satisfies $\varphi(x_1, \dots, x_r)$ in \mathbf{A} . The *atomic type* $\text{atyp}(\bar{a}, \mathbf{A})$ of a tuple $\bar{a} = (a_1, \dots, a_r) \in A^r$ is the set of atomic τ -formulas $\varphi(x_1, \dots, x_r)$ such that $\mathbf{A} \models \varphi(a_1, \dots, a_r)$. If \mathbf{A} is clear from context, we write $\text{atyp}(\bar{a})$ instead $\text{atyp}(\bar{a}, \mathbf{A})$.

We view graphs as $\{E\}$ -structures \mathbf{G} for E a binary relation symbol such that $E^{\mathbf{G}}$ is irreflexive and symmetric. Elements of G are *vertices*, elements of $E^{\mathbf{G}}$ are *edges*. A *subgraph* of \mathbf{G} is a graph $\mathbf{H} = (H, E^{\mathbf{H}})$ with $H \subseteq G$ and $E^{\mathbf{H}} \subseteq E^{\mathbf{G}}$. The *Gaifman graph* of a τ -structure \mathbf{A} is the graph $\mathbf{G}(\mathbf{A})$ with the same universe A as \mathbf{A} and $(a, a') \in E^{\mathbf{G}(\mathbf{A})}$ if a, a' are distinct and appear together in some tuple $\bar{a} \in R^{\mathbf{A}}$ for some relation symbol $R \in \tau$.

Let \mathbf{A}, \mathbf{B} be τ -structures. A *homomorphism from \mathbf{A} into \mathbf{B}* is a function $h: A \rightarrow B$ (where A, B are the universes of \mathbf{A}, \mathbf{B} respectively) such that $h(\bar{a}) \in R^{\mathbf{B}}$ for every relation symbol $R \in \tau$ and $\bar{a} \in R^{\mathbf{A}}$; here for an r -tuple $\bar{a} = (a_1, \dots, a_r) \in A^r$ we write $h(\bar{a})$ for the r -tuple $(h(a_1), \dots, h(a_r)) \in B^r$. Note $h: A \rightarrow B$ is a homomorphism from \mathbf{A} into \mathbf{B} if and only if for every tuple \bar{a} from A we have $\text{atyp}(\bar{a}, \mathbf{A}) \subseteq \text{atyp}(h(\bar{a}), \mathbf{B})$. Injective homomorphisms are *embeddings*. An *endomorphism of \mathbf{A}* is a homomorphism from \mathbf{A} to \mathbf{A} . An endomorphism of \mathbf{A} is *trivial* if it is the identity on A .

Throughout we mainly stick to the following notational conventions. Classes of structures are denoted by calligraphic letters, structures by boldface letters and their universes by the corresponding italic letter.

Path structures A *path* is a graph \mathbf{G} isomorphic to $([k], \{(i, j) \mid |i - j| = 1\})$ where $k := |G|$. Here, we write $[k] = \{1, \dots, k\}$ for $k \in \mathbb{N}, k \geq 1$. A sequence g_1, \dots, g_k such that $g_i \mapsto i$ is such an isomorphism is an *enumeration of \mathbf{G}* . The vertices g_1 and g_k are *endpoints*. The path is said to *connect* its endpoints and have *length* $k - 1$ (number of edges). If \mathbf{G} is a graph, then a *path in \mathbf{G}* is a subgraph of \mathbf{G} that is a path.

A *path structure (of vocabulary τ) \mathbf{P}* is a τ -structure whose Gaifman graph $\mathbf{G}(\mathbf{P})$ is a path. An *enumeration of \mathbf{P}* is an enumeration of $\mathbf{G}(\mathbf{P})$, and by an *endpoint of \mathbf{P}* we mean one of $\mathbf{G}(\mathbf{P})$. Note that a path structure has exactly two enumerations. A path structure \mathbf{P} (of vocabulary τ) is *rooted* if τ contains the unary relation symbol *root* such that $\text{root}^{\mathbf{P}}$ is a singleton containing one of the endpoints of $\mathbf{G}(\mathbf{P})$. For a rooted path structure \mathbf{P} with $|P| = k$, by an enumeration of \mathbf{P} we mean an enumeration p_1, \dots, p_k with $\text{root}^{\mathbf{P}} = \{p_1\}$. Note that a rooted path structure has exactly one enumeration. For $i \in [k - 1]$, we write e_i for the pair (p_i, p_{i+1}) ; we refer to the e_i as the *edges* of \mathbf{P} .

2.2 Parameterized logarithmic space

We consider (classical) problems Q as subsets of $\{0, 1\}^*$, the set of binary strings. Our model of computation are Turing machines with a read-only input tape, several work-tapes and a write-only output tape (the head does not move left and writes only 0, 1 and no blank). A Turing machine with oracle Q additionally has a write-only oracle tape special states “yes”, “no” and “?”; upon entering “?” the content y of the oracle tape is erased, its head placed on the first cell and state “yes” or “no” is entered according to whether $y \in Q$ or not; in such a step the machine is said to *query* y .

We follow [7] notationally. We view *parameterized problems* as pairs (Q, κ) where Q is a classical problem and $\kappa: \{0, 1\}^* \rightarrow \mathbb{N}$ is a *parameterization*. We assume that parameterizations are computable in logarithmic space, that is, the binary representation of $\kappa(x)$ is computable from x in space $O(\log |x|)$; here, $|x|$ is the length of $x \in \{0, 1\}^*$.

We exemplify how we present parameterized problems. The parameterized embedding problem associated with a class of structures \mathcal{A} is

$p\text{-EMB}(\mathcal{A})$
<i>Instance:</i> A structure $\mathbf{A} \in \mathcal{A}$ and a structure \mathbf{B} .
<i>Parameter:</i> $ A $.
<i>Problem:</i> Is there an embedding from \mathbf{A} into \mathbf{B} ?

The underlying classical problem is $\text{EMB}(\mathcal{A})$, the parameterization maps (\mathbf{A}, \mathbf{B}) to $|A|$.

The class para-L consists of those parameterized problems (Q, κ) decidable in *parameterized logarithmic space (with respect to κ)*, that is, space $f(\kappa(x)) + O(\log |x|)$ for some computable $f: \mathbb{N} \rightarrow \mathbb{N}$. Functions computable within this space are *pl-computable (with respect to κ)*. The class XL consists of those (Q, κ) decidable in space $f(\kappa(x)) \cdot \log |x|$ for some computable $f: \mathbb{N} \rightarrow \mathbb{N}$. These notions are from [2], our notation follows [6]. The class PATH has been introduced in [5]. It contains those (Q, κ) that are accepted by some nondeterministic algorithm (i.e. Turing machine) running in parameterized logarithmic space and which on input x makes at most $f(\kappa(x)) \cdot \log |x|$ many nondeterministic steps.

Theorem 2.1 ([3]) *Let \mathcal{P} be a decidable class of path structures. Then $p\text{-EMB}(\mathcal{P}) \in \text{PATH}$.*

A *pl-reduction* from (Q, κ) to (Q', κ') is a reduction R from Q to Q' which is pl-computable with respect to κ and such that there is a computable $f: \mathbb{N} \rightarrow \mathbb{N}$ such that $\kappa'(R(x)) \leq f(\kappa(x))$ for all $x \in \{0, 1\}^*$. A *pl-Turing reduction* from (Q, κ) to (Q', κ') is an algorithm with oracle Q' that decides Q , runs in parameterized logarithmic space with respect to κ , and has *bounded oracle access*: there is a computable $f: \mathbb{N} \rightarrow \mathbb{N}$ such that on input x the algorithm only queries y with $\kappa'(y) \leq f(\kappa(x))$.

Remark 2.2 Let \mathcal{P} be a decidable class of path structures of vocabulary τ . Any tuple in any relation in a structure in \mathcal{P} can have at most 2 distinct components. One can use this observation to give a pl-reduction of $p\text{-EMB}(\mathcal{P})$ to $p\text{-EMB}(\mathcal{P}')$ where \mathcal{P}' is a decidable class of “edge-coloured graphs”, i.e. path structures of some vocabulary τ' all of whose relation symbols have arity 2.

The following goes back to [5], in the form stated it appears in [3, Theorem 4.7].

Theorem 2.3 *The following is PATH-complete (with respect to pl-reductions).*

$p\text{-USTCON}$
<i>Instance:</i> A graph \mathbf{G} , vertices $s, t \in G$ and $\ell \in \mathbb{N}$.
<i>Parameter:</i> ℓ .
<i>Problem:</i> Is there a path connecting s and t in \mathbf{G} of length at most ℓ ?

This is a parameterized version of the classical *undirected s-t-connectivity problem* USTCON.

3 Main theorem statement

Given an ordered pair $e = (a, b)$ we will use the following non-standard notation. When $\ell \in \mathbb{N}$ we define $e^{-\ell}$ to be (a, b) when ℓ is even, and to be (b, a) when ℓ is odd.

Let \mathbf{P} be a rooted path structure with enumeration p_1, \dots, p_k where $k := |P|$. For $d \in \mathbb{N}$, we say e_i is *unfoldable of degree d* if $i > d$ and, for each $\ell \in [d]$, it holds that $\text{atyp}(e_i) \not\subseteq \text{atyp}(e_{i-\ell}^{-\ell})$. Note that all edges are unfoldable of degree 0, and being unfoldable of degree d implies being unfoldable of degree d'

for all $d' \leq d$. By *unfoldable* edges we mean edges unfoldable of degree 1. The *unfoldability degree* of \mathbf{P} is the sum $\sum_{i \in [k-1]} d_i$ where d_i is the maximal number such that e_i is unfoldable of degree d_i .

A class of rooted path structures \mathcal{P} has *bounded unfoldability degree* if there is a constant $c \in \mathbb{N}$ such that every $\mathbf{P} \in \mathcal{P}$ has unfoldability degree at most c .

Examples 3.1 Consider the path structures pictured in the Introduction. In the structures pictured in (1) and (4), no edge is unfoldable. In the structures pictured in (2), exactly the edges at even positions are unfoldable; they are unfoldable of degree 1 but not of degree 2. In each structure pictured in (3), only the last edge is unfoldable, of degree equal to the number of edges in the structure.

Theorem 3.2 (Main) *Let \mathcal{P} be a decidable class of rooted path structures of some vocabulary τ . If \mathcal{P} has bounded unfoldability degree, then there is a pl-Turing reduction of $p\text{-EMB}(\mathcal{P})$ to the parameterized problem*

p-LONGSHORT

Instance: A graph \mathbf{G} , vertices $s, t \in G$ and $k, \ell \in \mathbb{N}$ with $k < \ell$.

Parameter: ℓ .

Problem: Is it true that \mathbf{G} contains a path of length at least ℓ with endpoint s or a path of length exactly k connecting s and t ?

Otherwise, $p\text{-EMB}(\mathcal{P})$ is PATH-complete.

We divide the somewhat lengthy proof into lemmas proved in the next section. Lemma 4.1 together with Theorems 2.1 and 2.3 implies the first statement of Theorem 3.2. The second is proved as Lemma 4.2.

4 Proof of main theorem

4.1 Hardness

Lemma 4.1 *Let \mathcal{P} be a decidable class of rooted path structures. Assume \mathcal{P} does not have bounded unfoldability degree. Then there exists a pl-reduction from $p\text{-USTCON}$ to $p\text{-EMB}(\mathcal{P})$.*

Proof. Given an instance (\mathbf{G}, s, t, ℓ) of USTCON these reductions will proceed in two stages according to the characterization of parameterized logarithmic space as so-called *logarithmic space after a pre-computation* [7]. In the so-called *pre-computation* stage, the input parameter ℓ is mapped by a computable function to a pair (\mathbf{P}, X) such that $\mathbf{P} \in \mathcal{P}$ and X is a set of ℓ many edges of \mathbf{P} . The space required can be bounded by a computable function of ℓ . In our case, this first computation exploits special properties of \mathcal{P} .

In the second stage, the output (\mathbf{P}, \mathbf{B}) , an instance of $p\text{-EMB}(\mathcal{P})$ is produced. This computation is a logarithmic space computation that takes as input (\mathbf{G}, s, t, ℓ) plus (\mathbf{P}, X) . In our case, this computation is not going to depend on special properties of \mathcal{P} .

Recall that P denotes the universe of \mathbf{P} and G the universe of \mathbf{G} . The τ -structure $\mathbf{B} = \mathbf{B}(\mathbf{G}, \mathbf{P}, X, s, t)$ has universe $B := G \times P$. Let π_1 and π_2 denote the projections mapping an ordered pair $(g, p) \in B$ to its first resp. second component. Roughly speaking, we construct \mathbf{B} in such a way that an embedding h from \mathbf{P} into \mathbf{B} such that $h^* := \pi_2 \circ h$ is the identity on P yields a path of length at most $|X| = \ell$ in \mathbf{G} , and also vice-versa (this is formalized in Claims 1 and 3 below). We do this by ensuring that, whenever such h maps an edge e of \mathbf{P} to $(b, b') \in B^2$ then the first components of (b, b') are either equal or transverse an edge of \mathbf{G} ; the latter is ensured to happen only if $e \in X$. Proving correctness of our reduction will then amount to selecting an appropriate X such that we can prove that h^* is the identity on P .


We need some notation. Let $k := |P|$ and let p_1, \dots, p_k be the enumeration of \mathbf{P} . Recall we write e_i for the edge (p_i, p_{i+1}) . Let $1 \leq i_1 < \dots < i_\ell < k$ be such that $X = \{e_{i_1}, \dots, e_{i_\ell}\}$.

The structure \mathbf{B} interprets the unary relation symbol $root \in \tau$ by $root^{\mathbf{B}} := \{(s, p_1)\}$. To define the interpretation $R^{\mathbf{B}}$ of an r -ary relation symbol $R \in \tau$ we describe an algorithm \mathbb{A} that given an r -tuple $((g_1, q_1), \dots, (g_r, q_r)) \in B^r$ along with $\mathbf{G}, \mathbf{P}, X, s, t$ decides whether $((g_1, q_1), \dots, (g_r, q_r)) \in R^{\mathbf{B}}$:


1. check that $(q_1, \dots, q_r) \in R^{\mathbf{P}}$;
2. compute $i \in [k-1]$ such that every $q_j, j \in [r]$, equals p_i or p_{i+1} ;
3. check that there are $g, g' \in G$ such that every $(g_j, q_j), j \in [r]$, equals (g, p_i) or (g', p_{i+1}) ;
4. if $i+1 = k$, then check that $g' = t$;
5. if $i = i_1$, then check that $g = s$;
6. if $e_i \notin X$, then check that $g = g'$;
7. if $e_i \in X$, then check that $g = g'$ or $(g, g') \in E^{\mathbf{G}}$;
8. accept.

We understand that the computation is aborted and the algorithm rejects in case one of the checks fails. In particular, line 2 is entered only in case $(q_1, \dots, q_r) \in R^{\mathbf{P}}$; then the Gaifman graph $\mathbf{G}(\mathbf{P})$ of \mathbf{P} contains an edge between any two different components q_j 's; since $\mathbf{G}(\mathbf{P})$ is a path, the i asked for in line 2 is well-defined.

The following figure illustrates the construction. Consider the rooted path structure \mathbf{P} with universe $P = \{p_1, \dots, p_8\}$ that interprets $root$ by p_1 (depicted by the filled node) and a binary relation symbol R by the depicted arrows; further consider a directed graph \mathbf{G} with vertices $G = \{s, t, g, g'\}$ and directed edges again depicted by arrows:



The unfoldable edges in \mathbf{P} are e_2, e_3, e_5, e_6, e_7 . The $\{root, R\}$ -structure $\mathbf{B} = \mathbf{B}(\mathbf{G}, \mathbf{P}, \{e_2, e_6\}, s, t)$ exemplifies the construction in case 1 below. It has universe $G \times P$ and looks as follows. We draw a matrix of nodes with rows indexed by G and columns indexed by P . The interpretation of $root$ is $\{(s, p_1)\}$ and indicated by a filled node. The interpretation of R is by the arrows depicted.



Continuing with the proof, we make four observations concerning the structure $\mathbf{B} = \mathbf{B}(\mathbf{G}, \mathbf{P}, X, s, t)$.

Claim 1: If $(\mathbf{G}, s, t, \ell) \in \text{USTCON}$, then $(\mathbf{P}, \mathbf{B}) \in \text{EMB}(\mathcal{P})$.

Proof of Claim 1. Assume that $s = g_1, \dots, g_{\ell'+1} = t$ is (an enumeration of) a length $\ell' \leq \ell$ path connecting s and t in \mathbf{G} . Define the sequence $\tilde{g}_1, \dots, \tilde{g}_k$ in G as follows. The first i_1 many members $\tilde{g}_1, \dots, \tilde{g}_{i_1}$ equal $g_1 = s$, the next $i_2 - i_1$ members $\tilde{g}_{i_1+1}, \dots, \tilde{g}_{i_2}$ equal g_2 , and so on, the last $k - i_{\ell'}$ many members $\tilde{g}_{i_{\ell'}+1}, \dots, \tilde{g}_k$ equal $g_{\ell'+1} = t$. Then

$$(s, p_1) = (\tilde{g}_1, p_1), (\tilde{g}_2, p_2), \dots, (\tilde{g}_k, p_k) = (t, p_k)$$

is an enumeration of a copy of \mathbf{P} in \mathbf{B} , i.e. $p_i \mapsto (\tilde{g}_i, p_i)$ is an embedding from \mathbf{P} into \mathbf{B} . \dashv

Recall $\mathbf{G}(\mathbf{B})$ denotes the Gaifman graph of \mathbf{B} .

Claim 2: Let $\text{refl}(E^{\mathbf{G}}) := E^{\mathbf{G}} \cup \{(g, g) \mid g \in G\}$ be the reflexive closure of $E^{\mathbf{G}}$. Then

$$\pi_1 \text{ is a homomorphism from } \mathbf{G}(\mathbf{B}) \text{ into } (G, \text{refl}(E^{\mathbf{G}})); \quad (1)$$

$$\text{for all } e \in E^{\mathbf{G}(\mathbf{B})} : \text{ if } \pi_1(e) \in E^{\mathbf{G}}, \text{ then } \pi_2(e) \in X \text{ or } \pi_2(e)^{-1} \in X; \quad (2)$$

$$\pi_2 \text{ is a homomorphism from } \mathbf{B} \text{ into } \mathbf{P}. \quad (3)$$

Proof of Claim 2. Let $e = ((g, p), (g', p')) \in E^{\mathbf{G}(\mathbf{B})}$. Then there is $R \in \tau$ and $\bar{b} \in R^{\mathbf{B}}$ such that $(g, p), (g', p')$ both appear in \bar{b} . Since \mathbb{A} accepts \bar{b} , by line 2 there is $i \in [k-1]$ such that $p = p_i, p' = p_{i+1}$ or vice-versa. As one of the checks in line 6 or 7 is carried out, we have $(g, g') \in \text{refl}(E^{\mathbf{G}})$. Further, in case $\pi_1(e) = (g, g') \in E^{\mathbf{G}}$ we have $g \neq g'$, so by line 6 we must then have $e_i \in X$, implying that $\pi_2(e) \in X$ or $\pi_2(e)^{-1} \in X$. This shows (1) and (2).

Statement (3) is clear: if $\bar{b} \in R^{\mathbf{B}}$, then \mathbb{A} accepts \bar{b} , so $\pi_2(\bar{b}) \in R^{\mathbf{P}}$ by line 1. \dashv

Claim 3: Assume h is an embedding from \mathbf{P} into \mathbf{B} . Then $h^* := \pi_2 \circ h$ is an endomorphism of \mathbf{P} ; if h^* is trivial, then $(\mathbf{G}, s, t, \ell) \in \text{USTCON}$.

Proof of Claim 3. The first statement follows from (3). Assume h^* is trivial, that is, $h^*(p_i) = p_i$ for all $i \in [k]$. For $i \in [k]$ let $g_i \in G$ be such that $h(p_i) = (g_i, p_i)$. By (1), g_1, \dots, g_k satisfies $\pi_1(h(e_i)) = (g_i, g_{i+1}) \in \text{refl}(E^{\mathbf{G}})$ for all $i \in [k-1]$. By (2), $(g_i, g_{i+1}) \in E^{\mathbf{G}}$ only if $\pi_2(h(e_i)) = e_i \in X$. Hence the sequence g_1, \dots, g_k witnesses that g_1 and g_k are connected by a path of length at most $|X| = \ell$ in \mathbf{G} .

We are left to show $g_1 = s$ and $g_k = t$. The former holds as $h(p_1) = (g_1, p_1) \in \text{root}^{\mathbf{B}} = \{(s, p_1)\}$. To see $g_k = t$ note $h(e_{k-1}) \in \mathbf{G}(\mathbf{B})$ since h is an embedding. Hence, there are $R \in \tau$ and $\bar{b} \in R^{\mathbf{B}}$ with $(g_{k-1}, p_{k-1}), (g_k, p_k)$ appearing in \bar{b} . Since \mathbb{A} accepts \bar{b} we have $g_k = t$ by line 4. \dashv

Note that an endomorphism of \mathbf{P} may fail to be an endomorphism of its Gaifman graph $\mathbf{G}(\mathbf{P})$, for example, \mathbf{P} could have a constant endomorphism mapping each point p_i to the root p_1 . However, for endomorphisms of the form as in Claim 3, this can not happen:

Claim 4: Assume h is an embedding from \mathbf{P} into \mathbf{B} . Then $h^* := \pi_2 \circ h$ is an endomorphism of $\mathbf{G}(\mathbf{P})$, that is, for all $i \in [k-1]$ there is $j \in [k-1]$ such that $h^*(e_i) \in \{e_j, e_j^{-1}\}$.

Proof of Claim 4. By definition of \mathbb{A} we have for all $g, g' \in G$ and $j, j' \in [k]$:

$$\text{if } ((g, p_j), (g', p_{j'})) \in \mathbf{G}(\mathbf{B}), \text{ then } |j - j'| = 1. \quad (4)$$

For $i \in [k-1]$ choose g, g', j, j' such that $h(e_i) = ((g, p_j), (g', p_{j'}))$. Note h is also an embedding from $\mathbf{G}(\mathbf{P})$ into $\mathbf{G}(\mathbf{B})$. Hence $h(e_i) \in \mathbf{G}(\mathbf{B})$, so $|j - j'| = 1$ by (4), i.e. $h^*(e_i) = (p_j, p_{j'}) \in \{e_j, e_j^{-1}\}$. \dashv

We now exhibit a pl-reduction from $p\text{-USTCON}$ to $p\text{-EMB}(\mathcal{P})$ assuming \mathcal{P} does not have bounded unfoldability degree. The assumption implies that \mathcal{P} has at least one of the following properties:

(Case 1) For every $\ell \in \mathbb{N}$ there exists $\mathbf{P} \in \mathcal{P}$ such that at least ℓ many edges of \mathbf{P} are unfoldable.

(Case 2) For every $\ell \in \mathbb{N}$ there exists $\mathbf{P} \in \mathcal{P}$ such that at least one edge of \mathbf{P} is unfoldable of degree ℓ .

We exhibit a reduction as desired in both cases.

Case 1 Given an instance (\mathbf{G}, s, t, ℓ) of $p\text{-USTCON}$ the reduction first computes $\mathbf{P} \in \mathcal{P}$ such that \mathbf{P} contains at least ℓ many unfoldable edges. This can be done by computably enumerating the decidable class \mathcal{P} and testing for each structure output by the enumeration whether it has at least ℓ many unfoldable edges or not; the first structure passing the test is \mathbf{P} .

Recall P denotes the universe of \mathbf{P} . Again we write $k := |P|$ and let p_1, \dots, p_k be the enumeration of \mathbf{P} . From the input \mathbf{P} alone, the reduction computes ℓ many unfoldable edges $e_{i_1+1}, \dots, e_{i_\ell+1}$ with $1 \leq i_1 < \dots < i_\ell < k$. Then it computes $X = \{e_{i_1}, \dots, e_{i_\ell}\}$, the set of edges immediately preceeding the ℓ unfoldable ones, and outputs (\mathbf{P}, \mathbf{B}) for $\mathbf{B} = \mathbf{B}(\mathbf{G}, \mathbf{P}, X, s, t)$.

We have to show that

$$(G, s, t, \ell) \in \text{USTCON} \iff (\mathbf{P}, \mathbf{B}) \in \text{EMB}(\mathcal{P}). \quad (5)$$

Then we are done: as has already been observed, the algorithm \mathbb{A} runs in logarithmic space, so our reduction $(G, s, t, \ell) \mapsto (\mathbf{P}, \mathbf{B})$ can be computed in space $f(\ell) + O(\log |G|)$ for some computable $f: \mathbb{N} \rightarrow \mathbb{N}$. The output parameter $k = |P|$ is bounded by (in fact, equal to) a computable function of the input parameter ℓ . Thus, $(G, s, t, \ell) \mapsto (\mathbf{P}, \mathbf{B})$ defines a pl-reduction from $p\text{-USTCON}$ to $p\text{-EMB}(\mathcal{P})$.

We verify (5). The forward direction follows from Claim 1. Conversely, let $h: P \rightarrow B$ be an embedding from \mathbf{P} into \mathbf{B} . By Claim 3 it suffices to show that the endomorphism $h^* := \pi_2 \circ h$ is trivial.

Suppose not and choose the minimal $i \in [k]$ such that $h^*(p_i) \neq p_i$. Since h^* is an endomorphism (Claim 3) and $\text{root}^{\mathbf{P}} = \{p_1\}$ we have $h^*(p_1) = p_1$, so $i > 1$. Since $h^*(e_1)$ is an edge of $\mathbf{G}(\mathbf{P})$ (Claim 4), we have $h^*(e_1) = e_1$, so $i > 2$. Since $h^*(p_{i-1}) = p_{i-1}$ and $h^*(e_{i-1})$ is an edge of $\mathbf{G}(\mathbf{P})$ (Claim 4), we have $h^*(e_{i-1}) = e_i$ or $h^*(e_{i-1}) = e_{i-2}^{-1}$. As $h^*(p_i) \neq p_i$, we have $h^*(e_{i-1}) = e_{i-2}^{-1}$, that is,

$$h(e_{i-1}) = ((g, p_{i-1}), (g', p_{i-2}))$$

for certain $g, g' \in G$. Further, $\text{atyp}(e_{i-1}) \subseteq \text{atyp}(e_{i-2}^{-1})$ since h^* is an endomorphism of \mathbf{P} . Thus, e_{i-1} is not unfoldable and therefore $e_{i-2} \notin X$. Since $e_{i-1} \in \mathbf{G}(\mathbf{P})$ there exists $R \in \tau$ and a tuple $\bar{q} \in R^{\mathbf{P}}$ such that both p_{i-1} and p_i appear in \bar{q} . The image $h(\bar{q})$ contains (g, p_{i-1}) and (g', p_{i-2}) . Since $h(\bar{q}) \in R^{\mathbf{B}}$ and $e_{i-2} \notin X$, algorithm \mathbb{A} accepts carrying out the check in line 6, so $g = g'$.

For the $g'' \in G$ such that $h(p_{i-2}) = (g'', p_{i-2})$, we have $h(e_{i-2}) = ((g'', p_{i-2}), (g, p_{i-1}))$. As above we see that also $(g'', p_{i-2}), (g, p_{i-1})$ appear in some tuple in some relation from \mathbf{B} , and hence $g'' = g$.

Thus, $h(p_i) = h(p_{i-2}) = (g, p_{i-2})$ and h is not injective, a contradiction.

Case 2 In this case the reduction maps an instance (\mathbf{G}, s, t, ℓ) to (\mathbf{P}, \mathbf{B}) for $\mathbf{B} = \mathbf{B}(\mathbf{G}, \mathbf{P}, X, s, t)$ where $\mathbf{P} \in \mathcal{P}$ has an edge that is unfoldable of degree ℓ and X is the set of ℓ many edges preceeding this edge. More precisely, for $k := |P|$ let p_1, \dots, p_k be the enumeration of \mathbf{P} and let $e_i = (p_i, p_{i+1})$ be unfoldable of degree ℓ ; then X is $\{e_{i-1}, \dots, e_{i-\ell}\}$. Such (\mathbf{P}, X) can be computed from ℓ . We have to show:

$$(G, s, t, \ell) \in \text{USTCON} \iff (\mathbf{P}, \mathbf{B}) \in \text{EMB}(\mathcal{P}). \quad (6)$$

By Claim 1 it suffices to prove the backward direction in (6). So assume h is an embedding from \mathbf{P} into \mathbf{B} . By Claim 3 it suffices to show that the endomorphism $h^* := \pi_2 \circ h$ is trivial.

We first show that $h^*(p_j) = p_j$ for all $j \in [i - \ell]$ and in fact $h(p_j) = (s, p_j)$. Since $h(p_1) \in \text{root}^{\mathbf{B}} = \{(s, p_1)\}$ this holds for $j = 1$. Inductively, assuming $1 \leq j < i - \ell$ and $h(p_j) = (s, p_j)$ we show $g = s, q = p_{j+1}$ for g, q such that $h(p_{j+1}) = (g, q)$.

Since h is an embedding, $h(e_j) \in \mathbf{G}(\mathbf{B})$, so $\pi_1(h(e_j)) \in \text{refl}(E^{\mathbf{G}})$ by (1). Neither (p_j, q) nor (q, p_j) are in X since $j < i - \ell$. By (2) we have $\pi_1(h(e_j)) \in \text{refl}(E^{\mathbf{G}}) \setminus E^{\mathbf{G}}$, i.e. $s = g$. Now, $h^*(e_j) = (p_j, q)$ is an edge of $\mathbf{G}(\mathbf{P})$ (Claim 4), so equals e_j or e_{j-1}^{-1} . In the first case, $q = p_{j+1}$ and we are done. The other case is impossible: it implies $h(p_{j+1}) = (s, p_{j-1}) = h(p_{j-1})$, contradicting the injectivity of h .

Each of the ℓ many points $p_{i-\ell+1}, \dots, p_i$ is mapped by h^* to some $p_{i-\ell+j}$ with $j \geq 1$. Otherwise, there is $j \in [\ell - 1]$ such that $h^*(p_{i-\ell+j}) = p_{i-\ell+1}$ and $h^*(p_{i-\ell+j+1}) = p_{i-\ell}$, that is, $h^*(e_{i-\ell+j}) = e_{i-\ell}^{-1}$. Write $h(e_{i-\ell+j}) = ((g, p_{i-\ell+1}), (g', p_{i-\ell}))$ for certain $g, g' \in G$. Since this is an edge in $\mathbf{G}(\mathbf{B})$ there are $R \in \tau$ and $\bar{b} \in R^{\mathbf{B}}$ such that $(g, p_{i-\ell+1}), (g', p_{i-\ell})$ appear in \bar{b} . Since \mathbb{A} accepts \bar{b} , it follows from line 5 that $g' = s$ (note $i - \ell$ is the smallest index of an edge in X). Then $h(p_{i-\ell+j+1}) = (s, p_{i-\ell}) = h(p_{i-\ell})$, contradicting the injectivity of h .

In particular, $h^*(p_{i-\ell+1}) = p_{i-\ell+1}$ as otherwise $h^*(e_{i-\ell}) = (p_{i-\ell}, p_{i-\ell+j})$ for some $j > 1$ would not be an edge of $\mathbf{G}(\mathbf{P})$ (contradicting Claim 4).

Observe that in the graph $\mathbf{G}(\mathbf{B})$ each vertex $(g, p_{i-\ell+j})$ with $j \geq 1$ has only neighbors in $G \times \{p_{i-\ell+j+1}, p_{i-\ell+j-1}\}$ (cf. (4)). This implies that in the length ℓ sequence of points $h^*(p_{i-\ell+1}), \dots, h^*(p_i)$, the index increases or decreases by 1 in every step. It implies further that for none of these points h^* changes the parity of the index – more precisely: if ν, μ are indices $> i - \ell$ and $\leq i$ such that $h^*(p_\mu) = p_\nu$ then the parities of μ and ν are equal. Indeed, this follows easily by induction on μ with $i - \ell < \mu \leq i$: for the base case $\mu = i - \ell + 1$ note $h^*(p_{i-\ell+1}) = p_{i-\ell+1}$; the induction step follows from the previous observation that if μ increases by 1, then the index of $h^*(p_\mu)$ increases or decreases by 1, so changes its parity.

We next show that the index increases by 1 in every step, that is, $h^*(p_{i-\ell+j}) = p_{i-\ell+j}$ for all $j \in [\ell]$. Otherwise, $h^*(p_i) = p_\nu$ for some $i - \ell < \nu < i$. Then $h^*(e_i)$ equals e_ν or $e_{\nu-1}^{-1}$. We show both is impossible. Since the parity of ν equals that of i we have that $d := i - \nu$ is even. Thus, by the assumption that e_i is unfoldable of degree ℓ ,

$$\begin{aligned} \text{atyp}(e_i) \not\subseteq \text{atyp}(e_{i-d}^{-d}) &= \text{atyp}(e_\nu), \text{ and} \\ \text{atyp}(e_i) \not\subseteq \text{atyp}(e_{i-(d+1)}^{-(d+1)}) &= \text{atyp}(e_{\nu-1}^{-1}). \end{aligned}$$

We now know $h^*(p_j) = p_j$ for all $j \in [i]$. Then $h^*(p_{i+1}) = p_{i+1}$: otherwise, by Claim 4, $h^*(p_{i+1}) = p_{i-1}$, so $h^*(e_i) = e_{i-1}^{-1}$ and $\text{atyp}(e_i) \subseteq \text{atyp}(e_{i-1}^{-1})$ contradicting the unfoldability of e_i . It now follows that $h^*(p_j) = p_j$ for all $i < j \leq k$, and indeed $h(p_j) = (g, p_j)$ for $g := \pi_1(h(p_{i+1}))$. This is verified similarly as for the points p_j with $j \in [i - \ell]$. \square

4.2 Upper bound

Lemma 4.2 *Let \mathcal{P} be a decidable class of rooted path structures. Suppose \mathcal{P} has bounded unfoldability degree. Then there is a pl-Turing reduction from $p\text{-EMB}(\mathcal{P})$ to $p\text{-LONGSHORT}$.*

The proof uses color-coding, namely, we shall rely on the following lemma [7, page 349].

Lemma 4.3 *For every sufficiently large $n \in \mathbb{N}$, it holds that for all $k \leq n$ and for every k -element subset X of $[n]$, there exists a prime $p < k^2 \cdot \log n$ and $q < p$ such that the function $h_{p,q} : [n] \rightarrow \{0, \dots, k^2 - 1\}$ given by $h_{p,q}(m) := (q \cdot m \bmod p) \bmod k^2$ is injective on X .*

Proof.(of Lemma 4.2) Choose constants $c, d \in \mathbb{N}$ such that every structure $\mathbf{P} \in \mathcal{P}$ has at most c many unfoldable edges, and whenever an edge thereof is unfoldable of degree g , it holds that $g \leq d$.

Call an edge of a rooted path structure \mathbf{P} *critical* if its atomic type is different from the atomic type of the inverse of the edge preceding it. More precisely, let $k := |P|$ and p_1, \dots, p_k be the enumeration of \mathbf{P} ; recall we write e_i for the edge (p_i, p_{i+1}) of \mathbf{P} . For $i \in [k]$ the edge e_i is *critical* if $i > 1$ and $\text{atyp}(e_i) \neq \text{atyp}(e_{i-1}^{-1})$. This means e_i is either unfoldable or $\text{atyp}(e_i) \subsetneq \text{atyp}(e_{i-1}^{-1})$. Observe that e_2 is critical since $p_3 \notin \text{root}^{\mathbf{P}}$, so $\text{atyp}(e_2) \neq \text{atyp}(e_1^{-1})$. Thus every rooted path structure of size at least 3 (i.e. $k \geq 3$) has at least one critical edge.

For example, consider again the rooted path structures depicted in the Introduction. The critical edges of the structures in (1), (2) and (3) are the unfoldable ones (cf. Examples 3.1). The critical edges of the structures in (4) are the second one (unfoldable) and the first non-symmetric one (not unfoldable). Note that the structures in (3) and (4) have both exactly 2 critical edges. We have already seen that the first class has an associated embedding problem which is PATH-hard, and shall now see that the second one has an associated embedding problem which is pl-Turing reducible to $p\text{-LONGSHORT}$.

For $C \in \mathbb{N}$ let $\mathcal{P}(C)$ be the class of all $\mathbf{P} \in \mathcal{P}$ having at most C many critical edges. For every $C \in \mathbb{N}$ we are going to define a pl-Turing reduction from $p\text{-EMB}(\mathcal{P}(C))$ to $p\text{-LONGSHORT}$.

This suffices to prove the lemma:

Claim 5: There exists $C \in \mathbb{N}$ such that $\mathcal{P} = \mathcal{P}(C)$.

Proof of Claim 5. Let t denote the number of atomic τ -formulas in two variables. Let $\mathbf{P} \in \mathcal{P}$ and let e be an unfoldable edge of \mathbf{P} or the last edge. Let e' be the unfoldable edge preceding e (according to the enumeration of \mathbf{P}); if there is no such unfoldable edge, let e' be the first edge. Then there are $|\text{atyp}(e')| \leq t$ many critical edges between e and e' . In total, \mathbf{P} has at most $C := c + t(c + 1)$ many critical edges. \dashv

Let $\mathbf{P} \in \mathcal{P}$ with $k := |P| > 2 + cd + d$ and enumeration p_1, \dots, p_k . Then there exists a with $1 < a < k - d$ such that none of e_{a+1}, \dots, e_{a+d} are unfoldable. Fix a to be the minimum such value. As \mathbf{P} has at most c unfoldable edges, we have

$$1 < a \leq 2 + cd. \quad (7)$$

The following claim explains our interest in the number a .

Claim 6: For all $i \leq k - a$ we have $\text{atyp}(e_{a+i}) \subseteq \text{atyp}(e_a^{-i})$.

Proof of Claim 6. Suppose not. Choose i minimal such that $\text{atyp}(e_{a+i}) \not\subseteq \text{atyp}(e_a^{-i})$. Then for all $j \in [i]$ we have $\text{atyp}(e_{a+i-j}) \subseteq \text{atyp}(e_a^{-(i-j)})$, which implies $\text{atyp}(e_{a+i-j}^{-j}) \subseteq \text{atyp}(e_a^{-i})$, and thus $\text{atyp}(e_{a+i}) \not\subseteq \text{atyp}(e_{a+i-j}^{-j})$. This means e_{a+i} is unfoldable of degree i , so $i \leq d$. This contradicts the choice of a . \dashv

Our algorithm $\mathbb{A}(C)$ is going to be recursive, and the depth of the recursion bounded by a constant. It recurses on the parts $\mathbf{P} \uparrow i$ and $\mathbf{P} \downarrow i$ of \mathbf{P} obtained by “cutting \mathbf{P} at point i .” These structures are defined for $i \in [k]$ as follows: $\mathbf{P} \downarrow i$ is the substructure of \mathbf{P} induced on $\{p_1, \dots, p_i\}$, and $\mathbf{P} \uparrow i$ is obtained from the substructure of \mathbf{P} induced on $\{p_i, \dots, p_k\}$ by declaring p_i the new root, i.e. interpreting $\text{root} \in \tau$ by $\{p_i\}$. Without loss of generality we can assume that if \mathcal{P} contains \mathbf{P} , then it also contains all these structures $\mathbf{P} \uparrow i, \mathbf{P} \downarrow i$. If this would not be the case, we could add all these structures to \mathcal{P} and observe that the resulting bigger class still satisfies the assumptions of the lemma.

Let (\mathbf{P}, \mathbf{B}) be an instance of $p\text{-EMB}(\mathcal{P}(C))$ and p_1, \dots, p_k be the enumeration of \mathbf{P} . If $k := |\mathbf{P}| \leq 2 + cd + d$ or if $n := |\mathbf{B}|$ is not sufficiently large in the sense of Lemma 4.3 or if $n \leq k$, then $\mathbb{A}(C)$ uses “brute force”, that is, it runs an XL algorithm for $p\text{-EMB}(\mathcal{P}(C))$. So assume $2 + cd + d < k = |\mathbf{P}| < |\mathbf{B}| = n$ and n is sufficiently large. Then the algorithm $\mathbb{A}(C)$ computes the number a . There are two cases.

Case $a > 2$. In this case, $\mathbb{A}(C)$ loops through all functions $h: \{p_1, \dots, p_a\} \rightarrow B$. For every such h the algorithm checks whether it is an embedding of $\mathbf{P} \downarrow a$ into \mathbf{B} . If so, $\mathbb{A}(C)$ recurses on the instance $(\mathbf{P} \uparrow a, \mathbf{B}_h)$ where \mathbf{B}_h is obtained from \mathbf{B} and h as follows: take the substructure induced in \mathbf{B} with universe $B \setminus \{h(p_1), \dots, h(p_{a-1})\}$ and change the interpretation of root to $\{h(p_a)\}$. If the recursive call returns accepting, then $\mathbb{A}(C)$ halts and accepts.

By (7), each h can be stored using $O(\log n)$ bits. Hence, in the current case \mathbb{A} uses $O(\log n)$ space plus the space required by the recursive calls.

Case $a = 2$. In this case we shall use the following construction. It is similar to a construction used in [4, Theorem 18]. We can assume that the τ -structure \mathbf{B} has universe $B = [n]$ for some $n > 2$ sufficiently large in the sense of Lemma 4.3. Using the notation from this lemma, set

$$F := \{g \circ h_{p,q} \mid g: \{0, \dots, (k-1)^2 - 1\} \rightarrow \{2, \dots, k\} \text{ and } q < p < (k-1)^2 \log n\}.$$

For $f \in F$ and $b_1 \in B$ define the following graph $\mathbf{G}(f, b_1)$. Its vertices are $B \setminus \{b_1\}$. Its set of edges is the symmetric closure of the set of all $(b, b') \in (B \setminus \{b_1\})^2$ such that

- $f(b) + 1 = f(b')$ and $\text{atyp}(e_2^{-f(b)}, \mathbf{P}) \subseteq \text{atyp}((b, b'), \mathbf{B})$; and
- if $f(b) = 2$ then $\text{atyp}(e_1, \mathbf{P}) \subseteq \text{atyp}((b_1, b), \mathbf{B})$.

Informally, the idea is as follows. The first condition means to put an edge between all strongest (cf. Claim 7) $\text{atyp}(e_2)$ -edges in \mathbf{B} , but only between vertices with neighboring f -colours; the second condition ensures that a vertex with colour 2 is isolated unless it is a $\text{atyp}(e_1)$ -successor of b_1 in \mathbf{B} .

The following two claims pinpoint the properties we need this construction to have.

Claim 7: Let $b_1 \in B$, $f \in F$ and $s \in f^{-1}(2)$. If $\mathbf{G}(f, b_1)$ contains a length $k - 2$ path with endpoint s , then there is an embedding from \mathbf{P} into \mathbf{B} .

Proof of Claim 7. Let $s = b_2, b_3, \dots, b_k$ enumerate a path in $\mathbf{G}(f, b_1)$. Clearly, $\text{atyp}(e_1, \mathbf{P}) \subseteq \text{atyp}((b_1, s), \mathbf{B})$. For $1 < j \leq k - 1$ we have $|f(b_j) - f(b_{j+1})| = 1$. Since $f(s) = f(b_2) = 2$ it follows that $f(b_j)$ has the same parity as j . Thus, if $f(b_j) < f(b_{j+1})$, then

$$\text{atyp}(e_j) \subseteq \text{atyp}(e_2^{-j}) = \text{atyp}(e_2^{-f(b_j)}) \subseteq \text{atyp}((b_j, b_{j+1}), \mathbf{B}),$$

where the first inclusion holds by Claim 6; if $f(b_{j+1}) < f(b_j)$, then similarly

$$\text{atyp}(e_j) \subseteq \text{atyp}(e_2^{-j}) = \text{atyp}((e_2^{-f(b_{j+1})})^{-1}) \subseteq \text{atyp}((b_{j+1}, b_j)^{-1}, \mathbf{B}) = \text{atyp}((b_j, b_{j+1}), \mathbf{B}).$$

It follows that $p_j \mapsto b_j$ defines an embedding of \mathbf{P} into \mathbf{B} . —

Claim 8: Let $b_1 \in B$. The following are equivalent.

- (a) There exists $f \in F$ such that $\mathbf{G}(f, b_1)$ contains a length $k - 2$ path with one endpoint in $f^{-1}(2)$.
- (b) There exists $g \in F$ such that $\mathbf{G}(g, b_1)$ contains a path with one endpoint in $g^{-1}(2)$ and the other endpoint in $g^{-1}(k)$.

Proof of Claim 8. (b) implies (a) because any path in $\mathbf{G}(g, b_1)$ connecting points in $g^{-1}(2)$ and $g^{-1}(k)$ has length at least $k - 2$. Conversely, suppose b_2, \dots, b_k enumerates a path in $\mathbf{G}(f, b_1)$. By Lemma 4.3 there exists $g \in F$ such that $g(b_j) = j$ for all $2 \leq j \leq k$. Then b_2, \dots, b_k enumerates a path in $\mathbf{G}(g, b_1)$ with $b_2 \in g^{-1}(2)$ and $b_k \in g^{-1}(k)$. —

We now describe how algorithm $\mathbb{A}(C)$ works in the current case $a = 2$. It first loops through all tuples (b_1, p, q, g, s, t) such that $b_1 \in B, q < p < (k - 1)^2 \log n, g: \{0, \dots, (k - 1)^2 - 1\} \rightarrow \{2, \dots, k\}$ and $s, t \in B \setminus \{b_1\}$. For each such tuple \mathbb{A} first checks that $g(h_{p,q}(s)) = 2$ and $g(h_{p,q}(t)) = k$. Second it checks whether in $\mathbf{G}(g \circ h_{p,q}, b_1)$ the vertex s is connected by some path to t . If one such check is positive, the algorithm stops and accepts.

It follows from Claims 7 and 8, that if \mathbb{A} accepts here then indeed $(\mathbf{P}, \mathbf{B}) \in p\text{-EMB}(\mathcal{P}(C))$. A tuple (b_1, p, q, g, s, t) can be stored using $O(k \log k + \log n)$ bits. And the graph $\mathbf{G}(g \circ h_{p,q}, b_1)$ is pl-computable from the input. The second check can be done using Reingold's logarithmic space algorithm [9]. We thus see that this loop can be implemented within the allowed space.

If this first loop did not cause $\mathbb{A}(C)$ to accept, then $\mathbb{A}(C)$ recurses as follows. It computes the index of the first critical edge after e_2 , i.e. it computes the minimal $2 < i \leq k - 1$ such that e_{i+1} is critical. If there is no critical edge after e_2 , then $\mathbb{A}(C)$ sets $i := k - 1$. In both cases we have for all $2 \leq j \leq i$:

$$\text{atyp}(e_j) = \text{atyp}(e_2^{-j}). \tag{8}$$

Then $\mathbb{A}(C)$ loops a second time through all tuples (b_1, p, q, g, s, t) as before. It first checks that $g(h_{p,q}(s)) = 2$ and $g(h_{p,q}(t)) = i$. Then it queries the oracle whether

$$(\mathbf{G}(g \circ h_{p,q}, b_1), s, t, i - 2, k - 2) \in \text{LONGSHORT}.$$

If the oracle answers “no”, the algorithm considers the next tuple. If the oracle answers “yes”, then \mathbb{A} recurses on $(\mathbf{P} \uparrow i, \mathbf{B}')$ where \mathbf{B}' is obtained as follows. Take the induced substructure of \mathbf{B} with universe $B' := \{b \in B \mid i < g \circ h_{p,q}(b) \leq k\} \cup \{t\}$ and change the interpretation of *root* to $\{t\}$. If the recursive call returns accepting, then $\mathbb{A}(C)$ halts and accepts.

Note the instance \mathbb{A} recurses to is pl-computable from the input. So also in this second loop, $\mathbb{A}(C)$ uses parameterized logarithmic space plus the space required by the recursive calls. We argue for correctness: if $\mathbb{A}(C)$ accepts here, then $(\mathbf{G}(g \circ h_{p,q}, b_1), s, t, i - 2, k - 2) \in \text{LONGSHORT}$. By the fact, that $\mathbb{A}(C)$ entered the second loop, it follows that the statement Claim 8 (b) is false. Then the statement Claim 8 (a) is false. This implies that $\mathbf{G}(g \circ h_{p,q}, b_1)$ does not contain a length $k - 2$ path with endpoint s . Therefore, $(\mathbf{G}(g \circ h_{p,q}, b_1), s, t, i - 2, k - 2) \in \text{LONGSHORT}$ implies that there exists a path of length exactly $i - 2$ from s to t in $\mathbf{G}(g \circ h_{p,q}, b_1)$.

Let $s = b_2, \dots, b_i = t$ enumerate such a path. Then

$$g(h_{p,q}(b_j)) = j \quad (9)$$

for all $2 \leq j \leq i$. Assuming inductively that the recursive call accepts correctly, we have an embedding h from $\mathbf{P} \uparrow i$ into \mathbf{B}' . Since $\text{root}^{\mathbf{B}'} = \{t\}$ we have $h(p_i) = t = b_i$. Hence, $g(h_{p,q}(h(p_j))) > i$ for all $i < j \leq k$, and hence $h(p_j) \neq b_{j'}$ for all $2 \leq j' < i < j \leq k$ (by (9)). It follows that the following function h' is injective: map p_j to b_j for $1 \leq j \leq i$, and map p_j to $h(p_j)$ for $i < j \leq k$. Moreover, h' is an embedding from \mathbf{P} into \mathbf{B} : for $2 \leq j < i$ we have $h'(e_j) = (b_j, b_{j+1})$ and

$$\text{atyp}(e_j) = \text{atyp}(e_2^{-j}) \subseteq \text{atyp}((b_j, b_{j+1}), \mathbf{B}),$$

where the equality follows from (8), and the inclusion from (b_j, b_{j+1}) being an edge in $\mathbf{G}(g \circ h_{p,q}, b_1)$ and (9). We leave it to the reader to check $\text{atyp}(e_j) \subseteq \text{atyp}(h'(e_j), \mathbf{B})$ if $j = 1$ or $i \leq j \leq k - 1$.

We have argued that in all cases when $\mathbb{A}(C)$ accepts it does correctly so. Conversely, it is routine to verify, using Lemma 4.3, that $\mathbb{A}(C)$ accepts if $(\mathbf{P}, \mathbf{B}) \in \text{EMB}(\mathcal{P}(C))$. We conclude that $\mathbb{A}(C)$ decides $p\text{-EMB}(\mathcal{P}(C))$. Obviously, the oracle access is bounded. We are thus left to analyze the space complexity of $\mathbb{A}(C)$. We already argued that \mathbb{A} uses parameterized logarithmic space plus the space needed for the recursive calls. It is thus sufficient to check that the depth of the recursion is bounded by a constant.

In all cases $\mathbb{A}(C)$ recurses on $\mathbf{P} \uparrow i$ for some $i > 2$. The number of critical edges of $\mathbf{P} \uparrow i$ is less than or equal to the number of critical edges in \mathbf{P} . The same holds for the number of unfoldable edges. Indeed, each critical edge of $\mathbf{P} \uparrow i$ besides possibly its second one, is also critical in \mathbf{P} . And clearly every edge unfoldable in $\mathbf{P} \uparrow i$ is also unfoldable in \mathbf{P} .

In Case $a > 2$, the algorithm recurses on $\mathbf{P} \uparrow a$. That $a > 2$ implies that the first $a - 1$ edges in \mathbf{P} contain at least one unfoldable edge. It follows that $\mathbf{P} \uparrow a$ has less unfoldable edges than \mathbf{P} .

In Case $a = 2$, $\mathbb{A}(C)$ recurses on $\mathbf{P} \uparrow i$ for a certain $i > 2$. If this i equals $k - 1$, then $\mathbf{P} \uparrow i$ has only one edge and the recursive call uses “brute force” without any further recursion. If $i < k - 1$, then e_{i+1} is a critical edge. This becomes the second edge in $\mathbf{P} \uparrow i$. Since the critical e_2 is not present in $\mathbf{P} \uparrow i$, this structure has strictly less critical edges than \mathbf{P} .

It follows that in each recursive call either the number of unfoldable edges drops or “brute force” is applied or the number of critical edges drops. We conclude that the recursion depth is at most $c + C + 1$. \square

5 Dichotomy for rooted oriented paths

Let τ be the vocabulary $\{\text{root}, E\}$ where root is a unary relation symbol and E is a binary relation symbol. Let us say that a structure \mathbf{P} over τ is a *rooted oriented path* if it is a rooted path structure with enumeration p_1, \dots, p_k such that, for each $i \in [k - 1]$, exactly one of the two pairs (p_i, p_{i+1}) , (p_{i+1}, p_i) is contained in $E^{\mathbf{P}}$; and, no other pairs are in $E^{\mathbf{P}}$, in particular \mathbf{P} contains no loops (i.e. $E^{\mathbf{P}}$ is irreflexive). For such a structure and $C \geq 1$, let us say that the structure has a *C-alternating tail* if, for each $i \geq C$, the edge e_i is foldable if it exists (that is, if $i + 1 \leq k$). This means, that the edges e_C, e_{C+1}, \dots alternate in direction. For example,



pictures a rooted oriented path with a 4-alternating tail.

We establish the following dichotomy theorem.

Theorem 5.1 *Let \mathcal{A} be a decidable class of rooted oriented paths. If there exists $C \geq 1$ such that each structure in \mathcal{A} has a C -alternating tail, then $p\text{-EMB}(\mathcal{A})$ is in para-L. Otherwise, $p\text{-EMB}(\mathcal{A})$ is PATH-complete.*

Proof. For the hardness result, assume that there exists no constant C with the described property. By Theorem 3.2 it suffices to show that for each $c \geq 1$ there exists $\mathbf{P} \in \mathcal{P}$ of unfoldability degree at least $c - 1$. Indeed, if $\mathbf{P} \in \mathcal{P}$ does not have a c -alternating tail, then the last unfoldable edge e_i in \mathbf{P} satisfies $i \geq c$. Let e_{i_1}, \dots, e_{i_r} with $i_1 < \dots < i_r = i$ list the unfoldable edges. Set $i_0 := 1$ and observe that e_{i_j} is unfoldable of degree $i_j - i_{j-1}$. It follows that the unfoldability degree of \mathbf{P} is at least $i - 1 \geq c - 1$.

Now assume that there exists a constant $C \geq 1$ such that each $\mathbf{P} \in \mathcal{A}$ has a C -alternating tail. We have to find an algorithm deciding $p\text{-EMB}(\mathcal{A})$ in parameterized logarithmic space. This algorithm is akin to the oracle algorithm constructed in the proof of Lemma 4.2.

Let (\mathbf{P}, \mathbf{B}) be an instance of $p\text{-EMB}(\mathcal{A})$ and p_1, \dots, p_k be the enumeration of \mathbf{P} . We assume $k > C$ and $B = [n]$ for some $n > 2$ sufficiently large in the sense of Lemma 4.3. Set

$$F' := \left\{ g \circ h_{p,q} \mid g: \{0, \dots, (k - C + 1)^2 - 1\} \rightarrow \{C, \dots, k\} \text{ and } q < p < (k - C + 1)^2 \log n \right\}.$$

In the following we understand that f ranges over F' and h ranges over the set of functions from $\{p_1, \dots, p_C\}$ to B . The graph $\mathbf{G}(f, h)$ has vertices $B \setminus \{h(p_1), \dots, h(p_{C-1})\}$. Its set of edges is the symmetric closure of the set of those $(b, b') \in E^{\mathbf{B}}$ which satisfy

- $|f(b) - f(b')| = 1$, $f(b) \geq C$, $f(b') \geq C$ and $(p_{f(b)}, p_{f(b')}) \in E^{\mathbf{P}}$;
- if $f(b) = C$ then $b = h(p_C)$;
- if $f(b') = C$ then $b' = h(p_C)$.

Then the following are equivalent.

- (a) There are f, h such that h is an embedding of $\mathbf{P} \downarrow C$ into \mathbf{B} and $\mathbf{G}(f, h)$ contains a path with one endpoint in $f^{-1}(C)$ and the other endpoint in $f^{-1}(k)$.
- (b) There are f, h such that h is an embedding of $\mathbf{P} \downarrow C$ into \mathbf{B} and $\mathbf{G}(f, h)$ contains a length $k - C$ path with one endpoint in $f^{-1}(C)$.
- (c) $(\mathbf{P}, \mathbf{B}) \in \text{EMB}(\mathcal{A})$.

Indeed, that (a) and (b) are equivalent is seen similarly as Claim 8. That (b) implies (c) is seen similarly as Claim 7. Finally, that (c) implies (a) is easy to see using Lemma 4.3.

Algorithm \mathbb{B} on (\mathbf{P}, \mathbf{B}) checks that $|P| > C$ and $|B| \geq 2$ is sufficiently large in the sense of Lemma 4.3. If this is not the case, \mathbb{B} runs an XL algorithm for $p\text{-EMB}(\mathcal{A})$. Otherwise \mathbb{B} checks statement (a). This can be implemented within the allowed space similarly as explained for the second loop of the algorithm constructed in the proof of Lemma 4.2. \square

6 The long-short path problem is unavoidable

We show here that, up to pl-reduction, there is an embedding problem equivalent to the long-short path problem.

Theorem 6.1 *Let τ be the vocabulary containing the unary relation symbol root and a binary relation symbol E . There exists a class \mathcal{A} of rooted path structures of vocabulary τ such that $p\text{-EMB}(\mathcal{A})$ and $p\text{-LONGSHORT}$ are interreducible, with respect to pl-Turing reductions.*

Proof. For $0 \leq k < \ell$, define $\mathbf{P}_{k,\ell}$ to be the rooted path structure (on vocabulary τ) with universe $\{p_1, \dots, p_{\ell+1}\}$ where $\text{root}^{\mathbf{P}_{k,\ell}} = \{p_1\}$ and $E^{\mathbf{P}_{k,\ell}}$ is the union of $\{(p_i, p_{i+1}), (p_{i+1}, p_i) \mid i \in [\ell] \setminus \{k+1\}\}$ with $\{(p_{k+1}, p_{k+2})\}$. This structure has at most one unfoldable edge, namely the edge (p_{k+2}, p_{k+3}) (if it exists); this edge is unfoldable of degree 1, but not of degree 2. Let \mathcal{A} be the class containing all such structures $\mathbf{P}_{k,\ell}$. It follows from Theorem 3.2 that the problem $p\text{-EMB}(\mathcal{A})$ reduces to $p\text{-LONGSHORT}$.

We thus establish that $p\text{-LONGSHORT}$ reduces to $p\text{-EMB}(\mathcal{A})$. Let $(\mathbf{G}, s, t, k, \ell)$ be an instance of the problem $p\text{-LONGSHORT}$. The reduction produces the instance $(\mathbf{P}_{k,\ell}, \mathbf{G}')$ where \mathbf{G}' is defined as follows:

$$\begin{aligned} G' &:= G \cup \{q_1, \dots, q_{\ell-k}\}, \\ \text{root}^{\mathbf{G}'} &:= \{s\}, \\ E^{\mathbf{G}'} &:= E^{\mathbf{G}} \cup \{(t, q_1)\} \cup \{(q_i, q_{i+1}), (q_{i+1}, q_i) \mid i \in [\ell - k - 1]\}. \end{aligned}$$

Suppose that the original instance is a yes instance of $p\text{-LONGSHORT}$. If \mathbf{G} contains a path of length at least ℓ with endpoint s , then the structure $\mathbf{P}_{k,\ell}$ admits an injective homomorphism to \mathbf{G}' , namely, by simply mapping the elements of $\mathbf{P}_{k,\ell}$ onto the path. If \mathbf{G} contains an s - t path of length exactly k , then there is also an injective homomorphism; namely, the elements p_1, \dots, p_{k+1} are mapped onto the s - t path, with p_1 mapped to s and p_{k+1} mapped to t , and the elements $p_{k+2}, \dots, p_{\ell+1}$ are mapped to $q_1, \dots, q_{\ell-k}$, respectively.

Suppose that the created instance is a yes instance of $p\text{-EMB}(\mathcal{A})$; let h be the injective homomorphism witnessing this. If q_1 is not in the image of h , then none of the points q_i are, and hence h is an injective homomorphism into \mathbf{G} , implying that \mathbf{G} has a path of length ℓ with endpoint s . If q_1 is in the image of h , then it must hold that an edge of $\mathbf{P}_{k,\ell}$ maps onto (t, q_1) , since any path from s to q_1 in \mathbf{G}' must touch t immediately prior to touching q_1 . But since $(q_1, t) \notin E^{\mathbf{G}'}$, the only edge of $\mathbf{P}_{k,\ell}$ that can map onto (t, q_1) is (p_{k+1}, p_{k+2}) , implying that the image of p_1, \dots, p_{k+1} under h yields an s - t path of length k in \mathbf{G} . \square

Acknowledgements We thank the anonymous referees for their careful reading of the manuscript. Following their suggestions we added a number of examples, comments and discussions. The first author was supported by the Spanish Project MINECO COMMAS TIN2013-46181-C2-R, Basque Project GIU15/30, and Basque Grant UFI11/45. The authors wish to thank Eric Allender, Yijia Chen, and Michael Elberfeld for useful comments and discussion.

References

- [1] N. Alon, R. Yuster, and U. Zwick. Color-coding. *Journal of the ACM* 42(4):844–856, 1995.
- [2] L. Cai, J. Chen, R. G. Downey, and M. R. Fellows. Advice classes of parameterized tractability. *Annals of Pure and Applied Logic* 84(1):119–138, 1997.
- [3] H. Chen and M. Müller. The fine classification of conjunctive queries and parameterized logarithmic space. *Transactions on Computation Theory* 7(2): Article No. 7, 2015.

- [4] Y. Chen and M. Müller. Bounded variable logic, parameterized logarithmic space, and Savitch's theorem. In *Mathematical Foundations of Computer Science 2014 - 39th International Symposium, MFCS 2014, Budapest, Hungary, August 25-29, 2014. Proceedings, Part I*, pages 183–195, 2014.
- [5] M. Elberfeld, C. Stockhusen, and T. Tantau. On the space and circuit complexity of parameterized problems: Classes and completeness. *Algorithmica* 71(3):661–701, 2015.
- [6] J. Flum and M. Grohe. Describing parameterized complexity classes. *Information and Computation* 187(2): 291–319, 2003.
- [7] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer, 2006.
- [8] B. Lin. The parameterized complexity of k -biclique. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 605–615, 2015.
- [9] O. Reingold. Undirected connectivity in log-space. *Journal of the ACM* 55(4): Article No. 17, 2008.